

Presentation of Classical Propositional Tableaux on Program Design Premises*

Juan Michelini and Álvaro Tasistro

Universidad ORT Uruguay

michelini@ort.edu.uy, tasistro@ort.edu.uy

Abstract

We propose a presentation of classical propositional tableaux elaborated by application of methods that are noteworthy in program design, namely program derivation with separation of concerns. We start by deriving from a straightforward specification an algorithm given as a set of recursive equations for computing all models of a finite set of formulae. Thereafter we discuss the employment of data structures, mainly with regard to an easily traceable manual execution of the algorithm. This leads to the kinds of trees given usually as constituting the tableaux. The whole development strives at avoiding gaps, both of logical and motivational nature.

1 Introduction

We teach a course *Logic for Computing* in a Software Engineering programme of studies. Prior to this, students have received courses in Calculus, Algebra and introductory Programming in Java, plus a course called *Foundations of Computing*, which introduces polymorphic, higher-order functions and inductive types with the fundamental methods of induction and recursion in their various forms. *Foundations of Computing* makes emphasis on a mathematical approach to Programming, specifically on correctness proofs. *Logic for Computing*, in turn, concerns itself essentially with the notion of *formal* proof.

It follows from the foregoing that we should be very much interested in making explicit *methods* of proof. By this we mean both general strategies for developing and fully understanding solutions to problems, as well as manners of presenting the corresponding proofs which convey natural, concise and complete justifications of their design. Now, as it turns out, we have observed that some methods that have arisen within what could be called the science of Programming can be employed for obtaining or conveniently presenting mathematical results. This is to our mind a fact to be most welcome, for it exposes a unity of method between Programming and Mathematics that cannot but bring about positive outcomes for both sides, at least in as much the learning and teaching aspects are concerned.

In this paper we present an example of the latter, concerning the presentation of the method of *tableaux*. This is a proof procedure for both propositional and predicate logic dating back to [1] and [2], and whose ultimate variant (termed *analytic* tableaux) has been introduced in [3]. Specifically, what we do is: (1) We derive the method as a set of equations—to be used as rewriting rules—from a straightforward specification, namely the one demanding the computation of the set of all models of the given set of formulas. (2) We discuss the design of data structures for actually effecting and keeping trace of the execution of the method, which leads to the sorts of trees that are called “the tableaux” in textbooks. The first part yields a compact proof of the correctness of the method, much simpler than the ones in textbooks. The second part introduces the convenient and classical notation and establishes its correctness relating it to the set of equations originally given by a simple

* This work was partially supported by ANII–Agencia Nacional de Investigación e Innovación, Uruguay.



© Juan Michelini and Álvaro Tasistro;
licensed under Creative Commons License CC-BY

4th International Conference on Tools for Teaching Logic.

Editors: M. Antonia Huertas, João Marcos, María Manzano, Sophie Pinchinat, François Schwarzentruber; pp. 137–146



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



inductive argument. As a whole, the process is one in which we repeatedly employ simple techniques of *program derivation* and *separation of concerns* to obtain a presentation and justification both modular and simpler of the method of tableaux.

The rest of the paper consists of a general background section whose contents is assumed to be taught priorly to the study of tableaux. In section 3 we present the derivation of the equational algorithm calculating the set of all models of given set of formula. In section 4 we discuss the data structures for tracing the execution of the algorithm, leading to the usual presentations of tableaux, after which we finish up with a general discussion. The presentation is to be read basically as a concise course handout, with some explicit considerations of logical or didactic nature.

2 Background

Syntax. It is enough to consider the set of connectives $\{\neg, \wedge\}$. Then the set of formulæ is defined as usual, starting out from a denumerable set \mathcal{V} of propositional letters p :

$$\alpha, \beta ::= p \mid \neg\alpha \mid \alpha \wedge \beta.$$

We use *signed* formulae $\sigma ::= S\alpha$ where $S ::= F \mid T$, as the forms of assertion or judgement.¹

Semantics. Interpretations belong in $\mathcal{I} = \mathcal{V} \rightarrow \text{Bool}$. The *semantic value* of each formula is defined as follows — let A be the set of formulae and $(!)$ and $(\&\&)$ denote respectively Boolean negation and conjunction:

$$\begin{aligned} \llbracket _ \rrbracket &:: A \rightarrow \mathcal{I} \rightarrow \text{Bool} \\ \llbracket p \rrbracket^i &= i p \\ \llbracket \neg\alpha \rrbracket^i &= !\llbracket \alpha \rrbracket^i \\ \llbracket \alpha \wedge \beta \rrbracket^i &= \llbracket \alpha \rrbracket^i \&\& \llbracket \beta \rrbracket^i. \end{aligned}$$

Using the former we now define *truth* of an assertion (signed formula) in an interpretation. Call \hat{S} the boolean value corresponding to sign S . Then $i \models S\alpha \equiv \llbracket \alpha \rrbracket^i = \hat{S}$, which reads: i is a *model* of $S\alpha$, and also: i *satisfies* $S\alpha$ or $S\alpha$ is *valid in* i . We shall consider *finite* sets Γ of signed formulae and define models thereof (i.e. $i \models \Gamma$) as the interpretations satisfying every formula of Γ .

Truth in an interpretation. It is generally interesting to develop a method for checking truth of signed formulae in an interpretation. If we start with the propositional letters, we get:

$$\begin{aligned} i \models Sp &\equiv (\text{model of signed formula}) \\ \llbracket p \rrbracket^i &= \hat{S} \equiv (\text{semantic function}) \\ ip &= \hat{S}. \end{aligned}$$

For the other cases we wish to obtain (structurally) recursive equations. As to negation, writing \bar{S} the sign opposite to S , we obtain:

$$\begin{aligned} i \models S(\neg\alpha) &\equiv (\text{model of signed formula}) \\ \llbracket \neg\alpha \rrbracket^i &= \hat{S} \equiv (\text{semantic function}) \\ !\llbracket \alpha \rrbracket^i &= \hat{S} \equiv (\text{negating both sides to isolate } \llbracket \alpha \rrbracket^i) \\ \llbracket \alpha \rrbracket^i &= !\hat{S} \equiv (\text{opposite sign}) \\ \llbracket \alpha \rrbracket^i &= \hat{\bar{S}} \equiv (\text{model of signed formula}) \\ i \models \bar{S}\alpha. \end{aligned}$$

Finally, turning to conjunction:

$$i \models S(\alpha \wedge \beta) \equiv (\text{model of signed formula})$$

¹ The use of signed formulae avoids privileging one boolean value over the other and consequently simplifies definitions.

$$\begin{aligned} \llbracket \alpha \wedge \beta \rrbracket^i &= \hat{S} \equiv (\text{semantic function}) \\ \llbracket \alpha \rrbracket^i \&\& \llbracket \beta \rrbracket^i &= \hat{S}, \end{aligned}$$

where we seem to get stuck. Indeed, to rewrite the left hand side requires to consider the definition of $(\&\&)$ and this is not uniform with respect to truth and falsity. Therefore we are led to try instead distinguishing the cases of S :

$$\begin{aligned} \text{Case } T: \\ i \models T(\alpha \wedge \beta) &\equiv (\text{calculation above}) \\ \llbracket \alpha \rrbracket^i \&\& \llbracket \beta \rrbracket^i &= \text{True} \equiv (\text{definition of } \&\&) \\ \llbracket \alpha \rrbracket^i = \text{True} \wedge^2 \llbracket \beta \rrbracket^i &= \text{True} \equiv (\text{satisfaction}) \\ i \models T\alpha \wedge i \models T\beta. \end{aligned}$$

$$\begin{aligned} \text{Case } F: \\ i \models F(\alpha \wedge \beta) &\equiv (\text{calculation above}) \\ \llbracket \alpha \rrbracket^i \&\& \llbracket \beta \rrbracket^i &= \text{False} \equiv (\text{property of } \&\&) \\ \llbracket \alpha \rrbracket^i = \text{False} \vee \llbracket \beta \rrbracket^i &= \text{False} \equiv (\text{satisfaction}) \\ i \models F\alpha \vee i \models F\beta. \end{aligned}$$

Ultimately, we arrive at the following characterisation of the satisfaction relation:

$$\begin{aligned} \text{Signed letter :} \quad i \models Sp &\equiv i \models \mathbf{p} = \hat{S} \\ \text{Signed negation :} \quad i \models S(\neg\alpha) &\equiv i \models \overline{S}\alpha \\ \text{True conjunction :} \quad i \models T(\alpha \wedge \beta) &\equiv i \models T\alpha \wedge i \models T\beta \\ \text{False conjunction :} \quad i \models F(\alpha \wedge \beta) &\equiv i \models F\alpha \vee i \models F\beta. \end{aligned}$$

3 The Set of All Models

We now set ourselves the problem of computing *all models* of any given finite set Γ of signed formulae. This is accomplished by the function

$$\begin{aligned} \mathcal{M} &:: \mathcal{P}_{\text{fin}}(\Sigma) \rightarrow \mathcal{P}(\mathcal{I}) \\ \mathcal{M}(\Gamma) &= \{i \in \mathcal{I} \mid i \models \Gamma\}, \end{aligned}$$

where Σ is the set of signed formulae, \mathcal{P} is the power set operator yielding the set of subsets of given set, and \mathcal{P}_{fin} does the latter for the finite subsets³. Now this straightforward definition presents the inconvenience that, as a method of computation, it obliges to construct all the interpretations and check each of them against the formulae in Γ . We are rather in the search of a syntactic procedure, i.e. one that applied exclusively to the formulae in Γ ends up arriving at the desired set of models. Let us then examine Γ .

First of all, Γ could be empty, which is indeed a plainly uninteresting case. Indeed, every interpretation trivially satisfies the empty set of formulae and so the result in such case is \mathcal{I} . So let us assume $\Gamma \neq \emptyset$. If this is the case, then we can pick any one of the formulae σ in Γ and write the latter in the form $\Delta \mid \sigma$, which means that $\Gamma = \Delta \cup \sigma$ and $\sigma \notin \Delta$. Given the former, we can now write:

$$\begin{aligned} \mathcal{M}(\Gamma) &= (\text{split } \Gamma) \\ \mathcal{M}(\Delta \mid \sigma) &= (\text{definition of } \mathcal{M}) \\ \{i \in \mathcal{I} \mid i \models \Delta \mid \sigma\} &= (\text{satisfaction of a set of formulae}) \end{aligned}$$

² We distinguish \wedge from $\mathbf{\wedge}$, the first being the conjunction in the object language, the second being the conjunction in the meta-language. Likewise with \vee and $\mathbf{\vee}$.

³ Our use of “computing” seems at first generous indeed, since we are setting ourselves to generating in general infinite sets of infinite objects. Consider for that matter the case $\Gamma = \{\mathbf{p}\}$. Then any interpretation assigning *True* to \mathbf{p} is a member of the answer set. We shall see later how to settle this issue in detail —the general idea is to give a finite sufficient characterisation of infinite sets of interpretations.

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models \sigma\}.$$

The only source of information in the latter expression is the analysis of the form of σ , and so we are led to an examination of cases, i.e. to considering:

$$\mathcal{M}(\Delta \mid S \mathbf{p}) = \{i \in \mathcal{I} \mid i \models \Delta \wedge i \models S \mathbf{p}\},$$

$$\mathcal{M}(\Delta \mid S(\neg\alpha)) = \{i \in \mathcal{I} \mid i \models \Delta \wedge i \models S(\neg\alpha)\},$$

$$\mathcal{M}(\Delta \mid S(\alpha \wedge \beta)) = \{i \in \mathcal{I} \mid i \models \Delta \wedge i \models S(\alpha \wedge \beta)\}.$$

We can make profit of this analysis by using the results obtained at the end of the preceding section for checking the truth of signed formulae in a given interpretation. As it happens, the first case is a bit discouraging, for the satisfiability condition $i \models S \mathbf{p}$ takes us to consider the value of \mathbf{p} in the given interpretation, i.e. a semantic rather than a syntactic move. But it pays off to insist. Negation gives the following:

$$\mathcal{M}(\Delta \mid S(\neg\alpha)) = \text{(inferred above)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models S(\neg\alpha)\} = \text{(satisfaction of signed negation)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models \overline{S}\alpha\} = \text{(satisfaction of a set of formulae)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta, \overline{S}\alpha\} = \text{(definition of } \mathcal{M})$$

$$\mathcal{M}(\Delta, \overline{S}\alpha),$$

where we have used $(,)$ instead of (\cup) for set union. Notice that it is indeed this operation and not the formerly used split $(|)$ which is to be employed in this case, for we do not now know whether the formula $\overline{S}\alpha$ belongs or not to Δ . The equation thus obtained, namely

$$\mathcal{M}(\Delta \mid S(\neg\alpha)) = \mathcal{M}(\Delta, \overline{S}\alpha),$$

is very convenient, for it rewrites the desired set of all models into an expression in which the overall complexity of the formulae has been strictly decreased. The same works for conjunction, whose results with respect to satisfiability can be used by distinguishing the two cases of the sign affecting it:

$$\mathcal{M}(\Delta \mid T(\alpha \wedge \beta)) = \text{(inferred above)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models T(\alpha \wedge \beta)\} = \text{(truth of conjunction)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models T\alpha \wedge i \models T\beta\} = \text{(satisfaction of a set of formulae)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta, T\alpha, T\beta\} = \text{(definition of } \mathcal{M})$$

$$\mathcal{M}(\Delta, T\alpha, T\beta).$$

On the other hand:

$$\mathcal{M}(\Delta \mid F(\alpha \wedge \beta)) = \text{(inferred above)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models F(\alpha \wedge \beta)\} = \text{(falsity of conjunction)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge (i \models F\alpha \vee i \models F\beta)\} = \text{(distributing } \wedge \text{ over } \vee)$$

$$\{i \in \mathcal{I} \mid (i \models \Delta \wedge i \models F\alpha) \vee (i \models \Delta \wedge i \models F\beta)\} = \text{(trading } \wedge \text{ for } \vee \text{ out of set comprehension)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta \wedge i \models F\alpha\} \cup \{i \in \mathcal{I} \mid i \models \Delta \wedge i \models F\beta\} = \text{(satisfaction of a set of formulae)}$$

$$\{i \in \mathcal{I} \mid i \models \Delta, F\alpha\} \cup \{i \in \mathcal{I} \mid i \models \Delta, F\beta\} = \text{(definition of } \mathcal{M})$$

$$\mathcal{M}(\Delta, F\alpha) \cup \mathcal{M}(\Delta, F\beta).$$

As a result we have so far obtained:

$$\mathcal{M}(\Delta \mid S(\neg\alpha)) = \mathcal{M}(\Delta, \overline{S}\alpha)$$

$$\mathcal{M}(\Delta \mid T(\alpha \wedge \beta)) = \mathcal{M}(\Delta, T\alpha, T\beta)$$

$$\mathcal{M}(\Delta \mid F(\alpha \wedge \beta)) = \mathcal{M}(\Delta, F\alpha) \cup \mathcal{M}(\Delta, F\beta),$$

where the case of a signed letter, i.e. a *literal*, could not be included. Now taking a look at the preceding equations for \mathcal{M} , we readily realize that the missing case is actually that of a set not containing any composite formulae, i.e. that of a set of literals. Such is the base case of our recursion, since this proceeds by decreasing the size of the formulae of the set being treated—and not the size of the set itself. Therefore it is natural to wonder whether the solution of such base case could actually be just immediate. This is indeed the case,

because there is a straightforward manner of converting a set Γ of literals into the set of all its models. There are two cases:

- ▷ Γ contains pairs of opposite literals. Then it is inconsistent and the set of its models is \emptyset .
- ▷ Otherwise the models of Γ are the interpretations that coincide with Γ at the letters mentioned in it.

Formally, call $\hat{\Gamma}$ the set of all models of the set Γ of literals. It is defined as follows:

$$\hat{\Gamma} = \begin{cases} \emptyset & \text{if } \{S\mathbf{p}, \bar{S}\mathbf{p}\} \subseteq \Gamma \text{ for some } \mathbf{p}, \\ \{i \in \mathcal{I} \mid (\forall S\mathbf{p} \in \Gamma) i\mathbf{p} = \hat{S}\} & \text{otherwise.} \end{cases}$$

Notice that the alternative is decidable and that in the second case the result is sufficiently characterised by the set Γ of literals and so we get a finite representation of it. We can then put together equations for actually computing \mathcal{M} :

$$\mathcal{M}(\Delta \mid T(\alpha \wedge \beta)) = \mathcal{M}(\Delta, T\alpha, T\beta)$$

$$\mathcal{M}(\Delta \mid F(\alpha \wedge \beta)) = \mathcal{M}(\Delta, F\alpha) \cup \mathcal{M}(\Delta, F\beta)$$

$$\mathcal{M}(\Delta \mid S(\neg\alpha)) = \mathcal{M}(\Delta, \bar{S}\alpha)$$

$$\mathcal{M}(\Gamma) = \hat{\Gamma} \text{ if } \Gamma \text{ is a set of literals.}$$

We claim that \mathcal{M} captures the essence of the method of tableaux, and the derivation carried out above gives actually a quite simple proof of its correctness. Nevertheless, its actual execution needs to employ some kind of data structure to record the successive transformations leading to the final result. That is what we turn now to examining.

4 Data Structures for the Tableaux

List of lists. If we ignored the second equation above we would be in presence of a tail-recursive algorithm, i.e. one whose execution could consist merely in successively rewriting the finite set of formulae at hand. We would then do simply with a list of formulae from which we would choose the next formula to be transformed. Now, consideration of the second equation does not in principle introduce any dramatic modification of this situation: it is enough that each application of the equation produces a split of the list from which the formula $F(\alpha \wedge \beta)$ is taken into two lists, each of it containing exactly one of the two formulae $F\alpha$ and $F\beta$ in place of the original one, without any further change.

We illustrate this by means of an example. Suppose we consider $\Gamma = [T(\mathbf{p} \wedge \neg\mathbf{q}), F(\mathbf{p} \wedge \mathbf{q})]$, which we already make into a *list* of formulae (as indicated by the use of the $[\dots]$ notation). Then we may choose say the second formula to proceed, which leads us to employ \mathcal{M} 's second equation splitting the original list into two, yielding $[T(\mathbf{p} \wedge \neg\mathbf{q}), F\mathbf{p}], [T(\mathbf{p} \wedge \neg\mathbf{q}), F\mathbf{q}]$. In the next step we may choose any one of the two occurrences of the only composite (i.e. non-literal) formula. Say we take the left one. The equation to employ is \mathcal{M} 's first, yielding $[T\mathbf{p}, T(\neg\mathbf{q}), F\mathbf{p}], [T(\mathbf{p} \wedge \neg\mathbf{q}), F\mathbf{q}]$. Of course we may do the same with the right occurrence of the formula just considered, arriving at $[T\mathbf{p}, T(\neg\mathbf{q}), F\mathbf{p}], [T\mathbf{p}, T(\neg\mathbf{q}), F\mathbf{q}]$. We now are left only with the two occurrences of $T(\neg\mathbf{q})$ to treat, which we must do in two steps using the third equation of \mathcal{M} . We write the final result at once: $[T\mathbf{p}, F\mathbf{q}, F\mathbf{p}], [T\mathbf{p}, F\mathbf{q}, F\mathbf{q}]$. Clearly the first set of literals is unsatisfiable which, by the way, we could have noticed some steps earlier, thereby obtaining a less expensive development. The second set is just $\{T\mathbf{p}, F\mathbf{q}\}$ and characterizes all the models of the original Γ .

There seem to be three inconveniences as to this execution. The first is that we have treated one and the same formula twice, and that on two different occasions. One readily realizes that the issue is avoidable if the use of the branching equation corresponding to a false conjunction is always subsequent to the use of every other (non-branching) equation formerly applicable. The second inconvenience is that we have rewritten many a formula that was without change. And, finally, the execution is awkwardly traceable —we have

namely indicated the successive steps taken by means of narrative text interspersed in the successive rewritings. The latter is of importance when we consider executions by hand — then a more formal and easily checkable notation would be most welcome by both students and teachers. We shall consider these two remaining issues in the next two subsections, beginning with the latter about an easily traceable notation.

Tree of lists. The straightforward manner of making executions like the former traceable and easily verifiable is just to record the application of each rule, including mention to the formula used. We should therefore begin by naming the equations of the algorithm, say $T\wedge$, $F\wedge$, \neg and $!$ in the order in which they are written above. The procedure leads to the deployment of a tree structure whose nodes are lists of formulae as in the preceding section, and whose internal nodes (not leaves) are decorated by labels as explained presently:

0. To begin with, we have only one item, namely the original list of formula. This is of course a tree with only one terminal node (leaf).
1. At each step we choose a composite formula within a leaf (call this leaf \mathcal{L}) and apply the corresponding rule as already explained. As a result one or two new lists of formulae are obtained, which are linked to \mathcal{L} , becoming successors of \mathcal{L} in the tree. At the same time we label \mathcal{L} with the name of the equation and the formula used.

The leaves of these trees coincide with the lists of formulae obtained by the procedure explained in the preceding paragraph — we have only added a tree structure on top of them for tracing their computation. Therefore, the set of models of the root of the tree obtains as the union of the sets of models of the leaves. Formally, this much becomes clear after the consideration that the union of the sets of models of the leaves, and therefore the invariant just mentioned, are indeed preserved by each application of one equation as described above. Therefore the correctness of the computation procedure using these trees follows by straightforward mathematical induction. The right formulation and proof of this result is left as exercise.

Notice that the preceding description amounts to inductively defining these trees as a family $\mathcal{T}(\Gamma)$ indexed by the finite sets Γ in a manner such that the constructors stand in correspondence with the equations as named above, in the following manner: to internal nodes, constructors $T\wedge$, $F\wedge$ and \neg are associated, corresponding to the equation used in each case. The leaves are the as yet untreated nodes or those already formed by literals only. In either case we associate to the leaf the constructor $!$. Unfortunately, we must skip a detailed explanation for reasons of space.

Tree of formulae. The repetition of possibly large lists of formulae along the trees as introduced in the preceding section can be avoided, e.g. by employing the procedure described in [3]. We describe these less expensive trees as follows. The general idea is to write at each node of the tree different from the root only the formulae originated by the use (decomposition) of another formula. The root of the tree will contain the originally given set (list) of formulae. With this information it is possible to compute the full trees of the preceding paragraph provided the used formulae are recorded at each step, i.e. at each node. Therefore, the correctness of the present method with improved trees will follow from the correctness of the prior method. Specifically, we define the improved trees as follows:

1. Each node will have associated an explicit set E and an implicit set Γ of formulae. E is to be written down explicitly, whereas Γ is to be computed when necessary.
2. For the root of the tree, both E and Γ coincide with the originally given set of formulae.
3. For the other nodes, E will consist of one or two formulae.
4. All internal (i.e. non-leaf) nodes will also have associated one formula, to be called the one *used* at the node.

We now indicate how to extend the tree down from a leaf:

1. A formula in Γ is chosen and written down at the node as its used formula.
2. Then one proceeds according to the form of the chosen formula:
 - a. In case it is of the form $\top(\alpha \wedge \beta)$ then the tree is extended with *one* child node. For this new node, $E = \{\top\alpha, \top\beta\}$.
 - b. In case it is of the form $\text{F}(\alpha \wedge \beta)$ then the tree is extended with *two* children nodes. One of them will have $E = \{\text{F}\alpha\}$, whereas the other one will have $E = \{\text{F}\beta\}$.
 - c. Finally, in case the chosen formula is of the form $\text{S}(\neg\alpha)$ then the tree is extended with *one* child node having $E = \{\overline{\text{S}}\alpha\}$.

For every case of newly created node, the set Γ is computed as follows: If Γ_0 is the implicit set of the parent node, then $\Gamma = (\Gamma_0 - \sigma) \cup E$, where $-$ denotes deletion of a member in a set.

Now, to each improved tree t with a non-leaf root which has associated explicit set E and implicit set Γ of formulae, as well as used formula σ , a full tree of type $\mathcal{T}(\Gamma)$ can be associated, whose constructor is the one corresponding to the form of σ , i.e. $\top\wedge$, $\text{F}\wedge$ or \neg , and its children trees are the ones (recursively) corresponding to the children trees of t . If otherwise t is just a leaf, then its corresponding full tree is $\text{l}(\Gamma)$, where Γ is the implicit set of formulae of the leaf in question. This correspondence gives already a method for using the improved trees in order to compute all the models of any given set of formulae. Nevertheless, the following result makes such process easier: The implicit set at each leaf is the union of the explicit sets at the branch ending up at the leaf in question, minus those formulae that have been used on that branch. Thereby one can determine when a branch is *completed*, which happens when the implicit set at the corresponding leaf is a set Γ of literals. Further, then $\hat{\Gamma}$ is the corresponding set of models, and one can then compute the set of models of the whole tree (i.e. of the originally given set of formulae) by taking the union of the sets at each leaf, just as with the full trees.

5 Conclusions

We have put forward a presentation of classical propositional tableaux elaborated by application of some principles that are noteworthy in program design. Foremost among those principles is the one of *separation of concerns*: We have namely started by deriving from a straightforward specification an algorithm given as a set of recursive equations for computing all models of a finite set of formulae. The correctness of the algorithm is brought about hand-in-hand with its derivation by means of a basic inductive argument whose cases are each solved by calculational reasoning yielding identities between sets of interpretations that need not the usual “ping-pong” (or direct-and-converse) argument.

Thereafter we discussed the employment of data structures, mainly with regard to a manual execution of the algorithm. A requirement of natural traceability and verification led us to the trees of sets or lists of formulae presented in [1, 4], the correctness of which is immediate after their derivation as traces of the employment of the original equations. A further improvement avoids repetition of unmodified formulae giving rise to the trees presented in [3], whose correctness is in turn guaranteed by showing that they carry the same information as the former trees.

Smullyan’s classical presentation [3] introduces instead the method as a proof procedure for establishing unsatisfiability of (finite) sets of (signed) formulae. The tableaux are given directly in the form of our improved trees of formulae. The proof of correctness is then as usual composed by two arguments, one of soundness and one of completeness, to the

effect that unsatisfiable sets give rise to *closed* tableaux, i.e. one in which every branch contains a contradiction and thus has no model. The proof of soundness is by a quite direct tree induction, whereas the proof of completeness involves showing that an open *completed* branch, i.e. one in which every formula has been fully decomposed, is a Hintikka set. Besides, Hintikka's lemma is proven, to the effect that every Hintikka set has a model.

In our experience, the use of the method as in the classical presentation leads students to the realisation that they either prove the given set of formulae inconsistent or can compute every counter-example (i.e. a sufficient characterisation thereof). Subsequently they tend to ask why we cannot establish such fact as a meta-theoretical result. Our presentation does precisely that —and the correctness of the method as a proof procedure follows as immediate corollary. The idea of computing all models of the given set of formula has led us to give an abstract formulation of the procedure. We then treat as a separate matter the question of the concrete trace of the manual execution of the method. As we have been able to check, this treatment provides the students with improved command over the method, i.e. they exercise a more sound domain over what they are doing and also over the various possible notations or manners of justification they can give thereof.

It could be argued that Smullyan's presentation and proof is scalable to infinite sets of formulae and first-order-logic, and therefore ask about such feature regarding our presentation. Concerning infinite sets of formulae, the first thing to say is that the validity of our equations is certainly not affected. Nevertheless, they cannot of course be interpreted anymore as an algorithm. Even if we assume as usual a principle of omniscience concerning the infinite sets, the method of choice of the formulae to be successively decomposed by application of the equations is essential for getting the right result. But, as is the case also with the classical presentation, there exist method of orderly choice that guarantee (under the omniscience principle) the computation of all models and thus the correctness of the method. Generalisation to first-order logic, on the other hand, requires to abandon the idea of "computing all models", replacing it by e.g. "determining whether the set of formulae is or not (un)satisfiable".

We conclude that our presentation may contribute in a better way to the achievement of *profficiency with understanding*, which is our main learning objective. Also it emphasizes design methodology, which we strive to do along and across the whole of the program of studies. It also could be argued that the method is tailored to just students of Computing Science or Software Engineering. We however believe that it can be taught also without much difficulty to Mathematics or Philosophy students and that the advantages we claim to obtain can also be appreciated in such cases. This, however, is yet to be checked out.

Finally, we should like to think of this work as one interpretation and case of the disclosing of the "doing" of Mathematics as advocated by Dijkstra [5]. We have tried to avoid all gaps of both mathematical and motivational nature. To our mind, this case is yet another sample of the unity of structure and method that mathematics and programming share⁴. Exploiting such unity should be fruitful for improving understanding and thus better helping learning.

References

- 1 Hintikka, K. J. J. *Form and content in quantification theory*. Acta Philosophica Fennica, 8, 7-55, 1955.
- 2 Beth, E. W. *The Foundations of Mathematics*. North Holland 1959.

⁴ And that at a deeper level shows up in the propositions-as-types principle.

- 3 Smullyan R. M. *First-Order Logic*. Dover, 1994. An unabridged, corrected republication of the work first published by Springer-Verlag, New York, 1968.
- 4 Ben-Ari, M. *Mathematical Logic for Computer Science*. Springer, 2012.
- 5 Dijkstra, E. W. *EWD 1059*. <http://www.cs.utexas.edu/users/EWD/index10xx.html>

